

Claims

- [c1] 1. A system for use in achieving distributed data storage over a computer network comprising:
- a storage server system comprising one or more storage servers that each comprise a data storage device and a network interface for communicating with an application client that will require data storage and a management storage server; and
 - a management storage server system comprising one or more management storage servers that each comprise a network interface for communicating with an application client that will require data storage and each of said one or more storage servers;
- wherein each of said management storage servers comprises a data storage configuration identifier whose value is indicative of an allocation of data storage within said storage server system at a point in time;
- wherein an allocation of data storage within said storage server system comprises defining one or more virtual volumes of data storage distributed over one or more of said storage servers;
- wherein each of said management storage servers is capable of providing a first value for said data storage con-

figuration identifier to an application client;
wherein each of said management storage servers is capable of providing a second value for said data storage configuration identifier to each of said storage servers after there is a change in the allocation of data storage within said storage server system;
wherein each of said storage servers comprises a comparator capable of: (a) comparing said first value for said data storage configuration identifier which is associated with a data storage related request received from an application client with said second value for said data storage configuration, and (b) ignoring said data storage related request if said first value is not equal to said second value.

[c2] 2. A system, as claimed in claim 1, further comprises:
a driver for associating with an operating system of an application client;
wherein said driver is capable of associating said first value for said data storage configuration identifier with a data storage related request that is to be transmitted to a storage server over a network.

[c3] 3. A system, as claimed in claim 1, wherein:
each of said one or more management storage servers comprises a configuration map that is capable of identifying said one or more storage servers of said storage

server system;

at least one of said one or more management storage servers comprises an interface for allowing a user to alter said configuration map to add a storage server to said storage server system or subtract a storage server from said storage server system.

[c4] 4. A system, as claimed in claim 1, wherein:
each of said one or more management storage servers comprises a configuration map that is capable of: (a) identifying said one or more storage servers of said storage server system; (b) identifying a virtual volume that extends over two or more of said storage servers; and (c) identifying partitions of said virtual volume between two or more of said storage servers; and
at least one of said one or more management storage servers comprises an interface for allowing a user to alter said configuration map to define said partitions of said virtual volume.

[c5] 5. A system, as claimed in claim 1, wherein:
each of said one or more management storage servers comprises a configuration map that is capable of: (a) identifying said one or more storage servers of said storage server system; and (b) identifying a virtual volume that extends over two or more of said storage servers; and

at least one of said one or more management storage servers comprises an interface for allowing a user to direct that data from an application client reside on a first storage server associated with said virtual volume and a copy of the data reside on a second storage server associated with said virtual volume.

[c6] 6. A system, as claimed in claim 1, wherein:
each of said one or more management storage servers comprises a configuration map that is capable of: (a) identifying said one or more storage servers of said storage server system; and (b) identifying virtual volumes that each extend over two or more of said storage servers; and
at least one of said one or more management storage servers comprises an interface for allowing a user to direct that data be migrated from a first virtual volume to a second virtual volume.

[c7] 7. A system, as claimed in claim 1, wherein:
each of said one or more management storage servers comprises a configuration map that is capable of: (a) identifying said one or more storage servers of said storage server system; and (b) identifying virtual volumes that each extend over two or more of said storage servers; and
at least one of said one or more management storage

servers comprises an interface for allowing a user to direct that data be migrated from a first virtual volume to a second virtual volume and preserved on said first virtual volume.

- [c8] 8. A system, as claimed in claim 1, wherein:
at least one of said management storage servers is capable of changing the value of said data storage configuration identifier in response to a communication from a distributed lock manager that a lock will be revoked from a parallel database server.
- [c9] 9. A system, as claimed in claim 1, wherein:
said value for said data storage configuration identifier is a time stamp.
- [c10] 10. A system, as claimed in claim 9, wherein:
said time stamp comprises a logical time stamp.
- [c11] 11. A system, as claimed in claim 1, wherein:
said data storage device comprises a disk drive.
- [c12] 12. A system, as claimed in claim 5, wherein:
at least one of said management storage servers comprises a monitor for detecting when a synchronization/de-synchronization device indicates that the data residing on the second storage server may not be a copy of the data residing on the first storage server.

- [c13] 13. A system, as claimed in claim 12, wherein:
said synchronization/de-synchronization device comprises a first bit mask associated with said first storage server and a second bit mask associated with said second storage server.
- [c14] 14. A system, as claimed in claim 6, further comprising:
a migration status mechanism that allows a determination to be made as to whether a portion of the data on said first virtual volume has been migrated to said second virtual volume.
- [c15] 15. A system, as claimed in claim 14, wherein:
said migration status mechanism comprises a second virtual volume map with one or more pages that each represent a particular allocation of data space on said second virtual volume and a first virtual volume map with one or more pages that each correspond with one of the one or more pages of said second virtual volume map.
- [c16] 16. A system, as claimed in claim 15, wherein:
said migration status mechanism comprises a marking device for indicating in said first virtual map that a page of data has been migrated from said first virtual volume to said second virtual volume.

[c17] 17. A system, as claimed in claim 7, further comprising:
a migration status mechanism that allows a determination to be made as to whether a portion of the data on said first virtual volume has been migrated to said second virtual volume.